Overview of the I-WAY: Wide Area Visual Supercomputing

Thomas A. DeFanti

Electronic Visualization Laboratory
University of Illinois at Chicago
Chicago, IL 60639 USA
tom@eecs.uic.edu

Ian Foster, Michael E. Papka, and Rick Stevens

Mathematics and Computer Science Division
Argonne National Laboratory
Argonne, IL 60439 USA
{foster, papka, stevens}@mcs.anl.gov

Tim Kuhfuss

Electronics and Computing Technologies
Argonne National Laboratory
Argonne, IL 60439 USA
kuhfuss@anl.gov

Abstract

This article discusses the I-WAY project and provides an overview of the articles in this issue of *IJSA*. The I-WAY is an experimental environment for building distributed virtual reality applications and for exploring issues of distributed wide area resource management and scheduling. The goal of the I-WAY project is to enable researchers use multiple internetworked supercomputers and advanced visualization systems to conduct very large-scale computations. By connecting a dozen ATM testbeds, seventeen supercomputer centers, five virtual reality research sites, and over sixty applications groups, the I-WAY project has created an extremely diverse wide area environment for exploring advanced applications. This environment has provided a glimpse of the future for advanced scientific and engineering computing.

1 A Model for Distributed Collaborative Computing

The I-WAY, or Information Wide Area Year, was a year-long effort to link existing national testbeds based on ATM (asynchronous transfer mode) to interconnect supercomputer centers, virtual reality (VR) research locations, and applications development sites. The I-WAY was successfully demonstrated at Supercomputing '95 and included over sixty distributed supercomputing applications that used a variety of supercomputing resources and VR display environments. These applications groups have been pioneering the development of interactive distributed supercomputing applications during the past four years.

A major goal of the I-WAY project was to continue to work with the applications community to push the state of the art for distributed supercomputing. A primary thrust was applications that would use more than one supercomputer and one or more VR devices and that would begin exploring collaborative technologies to build shared virtual spaces in which to conduct computational science.

A related goal for the I-WAY project was to uncover the problems that are preventing widespread use of distributed supercomputing over ATM networks. Areas that were investigated included security, uniform computing environments, wide area scheduling and resource reservation, and distributed collaborative VR.

Future I-WAY research is focused on (1) making the I-WAY persistent, (2) understanding the performance constraints and requirement of applications, and (3) improving networking performance as seen by user applications codes.

Highlights of the I-WAY project include the following:

- Provided the first large-scale applications-driven test of IP/ATM for wide area supercomputing.
- Pushed the envelope on common computing tools across a wide variety of large-scale computers and VR environments.
- Motivated applications groups to consider using VR and supercomputing together.
- Established a framework for building a national research infrastructure in partnership with carriers, vendors, and the applications community.
- Built the first nationwide infrastructure to support collaborative computational science projects on a large scale.
- Developed a prototype software environment and wide area scheduler to enable easy use of distributed resources without detailed knowledge of their location and configuration, and with a high degree of security.

As a community-led effort, the I-WAY relied on the volunteer efforts of many people and groups to build and manage the network. The initial success of the project provides an excellent model for future advanced applications-driven networking.

2 I-WAY Networking

The principal goal of the networking group was to devise a scheme for connecting approximately twenty supercomputing centers over a wide area by using ATM (OC-3) networking technology. ATM was chosen rather than traditional Internet connectivity because it provides a broader bandwidth and is able to handle audio, video, and data more efficiently. The approach of ATM is somewhat like that of a phone call (Comer, 1995): specifically, a virtual circuit is established through the network before data is allowed to transfer. ATM also operates on fixed-length cells (53 bytes in length, with a 5-byte header and a 48-byte payload) (Vetter, 1995).

By using ATM Adaptation Layer 5 (AAL5), one of ATM's five adaptation levels, packets were sent across the I-WAY (Partridge, 1994). AAL5 enabled researchers to use standard TCP/IP protocols and tools. The I-WAY also supported direct ATM-oriented protocols for its users.

Much of the I-WAY's physical networking made use of existing smaller ATM research networks. These networks included the following:

AAInet - ACTS ATM Internetwork

ACTS - Advanced Communications Technology Satellite

ATDnet - Advanced Technology Demonstration Network

CalREN - California Research and Education Network

Canarie - Canadian Network for the Advancement of Research Industry and Education

CASA - Gigabit Testbed

DREN - Defense Research and Engineering Network

ESnet - Energy Sciences Network

MAGIC - Multidimensional Applications and Gigabit Internetwork Consortium

MREN - Metropolitan Research and Education Network

vBNS - Very High Performance Backbone Network Service

The separate networks were linked with the help of several major network service providers, including MCI, AT&T, Sprint, Ameritech, and Pacific Bell. Because of the combination of mixed hardware and the experimental nature of the I-WAY, permanent virtual paths (PVPs) were built to join the networks, and then permanent virtual circuits (PVCs) were constructed and entered into the routing tables by hand.

During the Supercomputing '95 conference a networking operations center was set up at San Diego that was used to interconnect the testbeds. Several OC-3 lines were brought into the conference center, including one from vBNS, one from AT&T, and two from Sprint.

Figure 1 presents a logical map of the I-WAY in December 1996.

3 Software Environment

A major challenge for the I-WAY was providing a uniform software environment across the geographically distributed and diverse computational resources. To meet this challenge, researchers developed a software infrastructure, called I-Soft, that provides a variety of services including scheduling, security (authentication and auditing), parallel programming support (process creation and communication), and a distributed file system.

The I-Soft system was designed to run on dedicated I-WAY point of presence (I-POP) machines (see Figure 2).

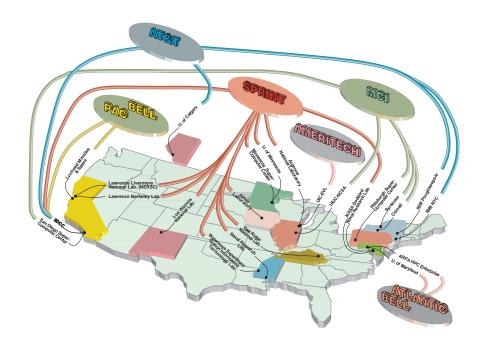


Figure 1: Logical Map of the I-WAY (Richard Foster and Linda Winkler, ANL, 1996)

Each I-POP is accessible via the Internet, but operates inside a site's firewall. An ATM interface allows it to monitor and, in principle, manage the site's ATM switch; it also allows the I-POP to use the ATM network for system management traffic. Site-specific implementations of a simple remote resource management interface allow I-POP systems to communicate with other machines at the site to allocate resources to users, start processes, and so forth. The Andrew distributed file system (AFS) is used as a repository for system software and status information.

A critical component of the software environment was the distributed scheduler. Political and technical constraints made it infeasible to provide a single "I-WAY scheduler" to replace the schedulers that are already in place at various sites. Instead, we implemented a two-part strategy that allowed administrators to configure dedicated resources into virtual machines and allowed users to request time on particular virtual machines. The strategy involved a (1) central scheduler daemon that managed and allocated time on the different virtual machines on a first-come, first-served basis, and (2) a local scheduler daemon that communicated directly with the local site scheduler. Local schedulers performed site-dependent actions in response to requests from the central scheduler to allocate resources, create processes, and deallocate resources (Foster et al., 1996b).

Security is a major and multifaceted issue in I-WAY-like systems. Authentication to I-POPs was handled by using a telnet client modified to use Kerberos authentication and encryption. This approach avoided the need for passing passwords in clear text over the network. The scheduler software kept track of which user id was to be used at each site for a particular I-WAY user and served as an "authentication proxy," performing subsequent authentication to other I-WAY resources on the user's behalf. This proxy service was invoked each time a user used the command language described above to allocate computational resources or to create processes.

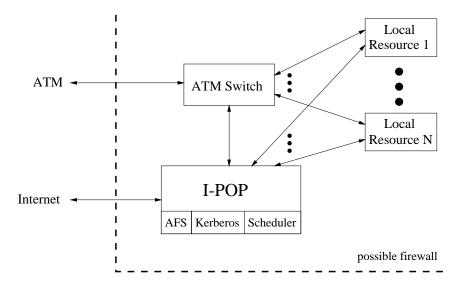


Figure 2: An I-WAY Point of Presence (I-POP) machine

A variety of parallel programming tools were provided on the I-WAY system. The challenge here was to enable users to exploit the features of the various computers and yet to hide the unnecessary details of networks and systems. To this end, we adapted the Nexus multithreaded communication library (Foster et al., 1994, 1996a) to execute in the I-WAY environment. Nexus supports automatic configuration mechanisms that allow it to use information contained in resource databases to determine which startup mechanisms, network interfaces, and protocols to use in different situations. Several other libraries, notably the CAVEcomm virtual reality library (Disz et al., 1995a, 1995b) and the MPICH (Gropp et al., 1996) implementation of MPI, were extended to use Nexus mechanisms.

4 Supercomputer Sites

Early in the I-WAY project all major U.S. supercomputer sites were invited to participate in the I-WAY. Of the approximately seventy sites, fewer than half were connected to an existing ATM testbed, and others were unavailable because of resource limitations or other constraints. The final set of I-WAY supercomputer resource sites numbered approximately twenty:

- DOE supercomputer centers
 - Argonne National Laboratory
 - Los Alamos National Laboratory
 - National Energy Research Supercomputer Center
 - Oak Ridge National Laboratory
 - Sandia National Laboratories
- National Science Foundation supercomputer centers
 - National Center for Supercomputing Applications

- Cornell
- San Diego Supercomputer Center
- Pittsburgh Supercomputing Center
- National Aeronautics and Space Administration supercomputers
 - Ames
 - Goddard
 - Caltech/JPL
- Lockheed Martin Missiles and Space
- Advanced Research Projects Agency, Enterprise Integration Testbed
- Maui High Performance Computing Center
- U.S. Army Waterways Experiment Station
- University of Illinois at Chicago, Electronic Visualization Laboratory

I-WAY testbeds were set up comprising a set of computing resources (processors, disk, network connections) chosen from the resources available at each site. These testbeds were "named" and made available from the scheduler as a virtual machine. The named virtual testbeds allowed users to think in terms of a wide area collection of nodes of a particular type. For example, a collection IBM SP2 nodes at Argonne and Cornell could make up a particular testbed and be scheduled as a unit.

The linked sites provided access to a wide variety of supercomputers, including the Thinking Machines CM-5, SGI Power Challenge Array, IBM SP2, Cray T3D and C90, and Intel Paragon. The computer power represented the largest ensemble of distributed computing resources ever interconnected on a high-performance wide area network.

5 Virtual Reality Environments

Visualizing data generated at remote supercomputer sites required high-bandwidth display devices. Researchers also wished to be able to interact with their supercomputing simulations in real time. The presentation and interaction with the simulation were handled by three different virtual reality display systems (Korab and Brown, 1995).

5.1 CAVE

The CAVETM (CAVE Automatic Virtual Environment) is a 10x10x9 foot room that makes use of rear-projected high-resolution projectors to produce an immersive 3D environment (Figure 3). The CAVE environment, originally developed by the Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago, produces a 3D stereo effect by displaying in alternating succession the left and right eye views of the scene as rendered from the viewers perspective (Cruz-Neira et al., 1993). These views are then seen by the user through a pair of

LCD shutter glasses whose lenses open and close forty-eight times a second in synchronization with the left and right eye views. The correct viewer centered projection is calculated based on the viewer's position and orientation as determined by a electromagnetic tracking system. The position and orientation of a 3D wand are also tracked; this wand allows for navigation of and input into the virtual world. Along with the visual feedback of the CAVE environment, a complete 3D audio environment is available to the user.

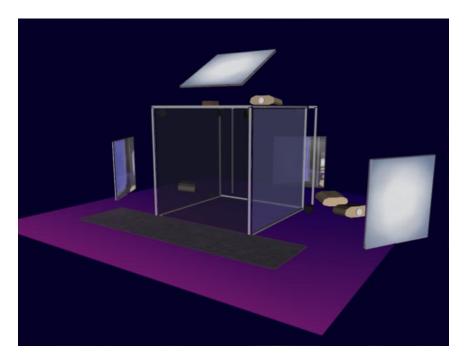


Figure 3: CAVE Virtual Environment (Milana Huang, EVL, 1994)

5.2 ImmersaDesk

Two ImmersaDeskTM systems were used for demonstrations at Supercomputing '95. An ImmersaDesk is based on the same rear-projection technology as the CAVE (Figure 4). It is a fully interactive, 3D immersive environment that is about the size of a large drafting table. The ImmersaDesk allows for one tracked viewer, along with two to three passive viewers.

5.3 NII/Wall

The third display technology used at Supercomputing '95 as a outlet for the applications was the NII/Wall (Figure 5). The NII/Wall is a large rear-projected system that is created from four 1280 by 1024 screens. The NII/Wall can be used as a large ImmersaDesk, where the images are projected in stereo and the viewer is tracked, or can substitute for a large high-resolution workstation. The NII/Wall was developed by EVL, the National Center for Supercomputing Applications, and the University of Minnesota, with support from Silicon Graphics, Inc.



Figure 4: ImmersaDesk Virtual Environment (Jason Leigh, EVL, 1995)

6 Applications

The I-WAY project was applications driven. By that we mean that it was intended as a resource for large-scale, high-performance applications requiring greater computer power and more advanced visualization environments than typically available at a single supercomputer site.

The applications were chosen from over one hundred projects in a national call for applications to use the I-WAY environment. The final applications fell into one of five classifications (Figure 6): Supercomputer - Supercomputer, Remote Resource - VR, VR - VR, Multisupercomputer - Multi-VR, or Video, WWW, GII-Window (DeFanti and Stevens, 1995).

The applications comprised nineteen different scientific and engineering disciplines:

Astronomy
Astrophysics
Atmospheric Science
Biochemistry, Molecular Biology, and Structural Biology
Biological and Medical Imaging
Chemistry
Distributed Computing
Earth Science
Education
Engineering
Geometric Modeling
Material Science

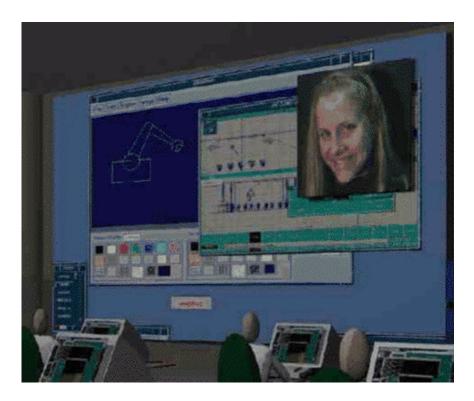


Figure 5: NII/Wall Virtual Environment (Jason Leigh, EVL, 1995)

Mathematics
Microphysics and Macrophysics
Neuroscience
Performance Analysis
Plasma Physics
Teleoperations/Telepresence
Visualization

This special issue of *IJSA* focuses on nine of these applications. Below, we briefly summarize each of the applications.

The article "Collaborative Virtual Environments Used in the Design of Pollution Control Systems" by Diachin et al. describes the development of a virtual reality environment for the interactive analysis and design of injective pollution control systems. Several mechanisms for the real-time computation and visualization of particle dynamics in fluid flow and spray simulations are described. The required computations are performed on an IBM SP2 processor and communicated to one or more CAVE Automatic Virtual Environments using the CAVEcomm message-passing library. The article analyzes both the computational and communication costs for typical interactions with the model and examines the network bandwidth required for real-time response.

The article "Performance Modeling of Interactive, Immersive Virtual Environments for Finite Element Simulations" by Taylor et al. analyzes the components of the lag time resulting from the execution of a finite element simulation on the IBM SP2 supercomputer at Argonne and displayed in an interactive virtual environment at Supercomputing '95 in San Diego. The article presents results comparing the I-WAY network with that of the traditional Internet. In

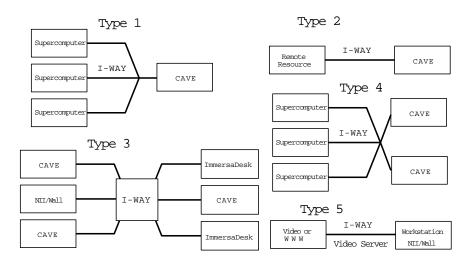


Figure 6: Five Application Types That Used The I-WAY

addition, the article discusses the components of end-to-end lag.

The article "The Chesapeake Bay Virtual Environment (CBVE): Initial Results from the Prototypical System" by Wheless et al. discusses the use of virtual environments to visualize ecological datasets. The datasets are generated by numerical supercomputer simulations that provide insight into the ecological impact of physical and biological interactions.

The article "Implementing a Collaboratory for Microscopic Digital Anatomy" by Young et al. introduces a system that will allow biologists around the world to acquire, over the network, data taken from an electron microscope. Using remote high-performance computers, biologists can process and enhance three-dimensional datasets and then view the results on their local display device.

The article "Interactive Scientific Exploration of Gyrofluid Tokamak Turbulence" by Kerbel et al. examines the use of an interactive visual environment coupled to high-performance computations. The article discusses previous experience with such coupling efforts and discusses the demonstration at Supercomputing '95. The system connects a Cray T3D running the physics simulations and postprocessing visualization code to an interactive virtual environment.

The article "Exploring Coupled Atmosphere-Ocean Models Using Vis5D" by Hibbard et al. discusses the use of a distributed client/server system for viewing large datasets. Using an IBM SP2 at Argonne National Laboratory as a server, researchers fed data across the I-WAY to a pair of SGI Onyx machines at Supercomputing '95. The SGI machines are used to process the graphics for display in the CAVE, where viewers are allowed to interact with the data.

The article "Radio Synthesis Imaging – A Grand Challenge HPCC Project" by Crutcher et al. discusses the issues of real-time transfer and archiving of data from remote telescopes, the processing of data on supercomputers for image improvement, and the final archiving of images into a digital library. In particular, the article discusses these issues as they applied to the demonstration at Supercomputing '95.

The article "Early Experiences with Distributed Supercomputing on I-WAY: First Principles Materials Science and Parallel Acoustic Wave Propagation" by Geist et al. introduces the details and results of their materials science simulation and seismic simulation. The article also describes the library used to steer and visualize both applications.

The article "Galaxies Collide on the I-WAY: An Example of Heterogeneous Wide-Area Collaborative Supercomputing" by Norman et al. describes the attempt to use the resources of the NSF supercomputing centers as a metacomputer. The article discusses the development of a simulation that is independent of the host architecture and that takes advantage of resources on the local compute servers.

7 Conclusions and Future Directions

The I-WAY project was a remarkable achievement. A community effort, it required unprecedented cooperation among telecommunication providers, equipment vendors, and applications scientists nationwide—most on a volunteer basis.

It also required meeting an extremely tight deadline. The project was initiated in April 1995; by May, the concept had been demonstrated with a transcontinental virtual underwater experiment; by fall, parallel languages, scalable Unix tools, graphics libraries, and communications software had been integrated into a cohesive environment; and by December 1995, more than sixty high-performance applications were successfully demonstrated at the Supercomputing '95 conference.

We are now working to address some of the critical research issues identified in I-WAY project. The Globus project (see http://www.mcs.anl.gov/globus/) is addressing issues of resource location (computational resource brokers), automatic configuration, scalable trust management, and high-performance distributed file systems. In addition, we and others are defining and constructing future I-WAY-like systems that will provide further opportunities to evaluate management and application programming systems.

Acknowledgments

The I-WAY was a multi-institutional, multi-individual effort. Maxine Brown, Linda Winkler, Mary Spada, and Remy Evard played major roles. We acknowledge Gail Pieper for all her help in preparing this paper. This work was supported in part by the Mathematical, Information, and Computational Sciences Division subprogram of the Office of Computational and Technology Research, U.S. Department of Energy, under Contract W-31-109-Eng-38. Support also provided by U.S. Department of Energy, and the National Science Foundation, grant CDA-9303433, with major support from the Advanced Research Projects Agency.

References

- [1] Comer, D. E. 1995. Internetworking with TCP/IP, 3rd. ed. Englewood Hall, Prentice Hall.
- [2] Cruz-Neira, C., Sandin, D. J., and DeFanti, T. A. 1993. Surround-screen projection-based virtual reality: The design and implementation of the CAVE. In *Computer graphics* (*Proceedings of SIGGRAPH '93*), ACM SIGGRAPH, pp. 135–142.
- [3] DeFanti, T., and Stevens, R. 1995. Notes from I-WAY workshop April 18–20, at Argonne National Laboratory. URL http://www.iway.org/.

- [4] Disz, T. L., Papka, M. E., and Pellegrino, M. 1995a. CAVEcomm programmers manual.
- [5] Disz, T. L., Papka, M. E., Pellegrino, M., and Stevens, R. 1995b. Sharing visualization experiences among remote virtual environments. In *International workshop on high performance computing for computer graphics and visualization*. Berlin, Springer-Verlag, pp. 217–237.
- [6] Foster, I., Kesselman, C., and Tuecke, S. 1994. The Nexus task-parallel runtime system. In *Proc. 1st Intl Workshop on Parallel Processing*, pages 457-462. Tata McGraw Hill.
- [7] Foster, I., Geisler, J., Kesselman, C., and Tuecke, S. 1996a. Multimethod communication for high-performance networked computing systems. Preprint, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Ill.
- [8] Foster, I., Geisler, J., Nickless, W., Smith, W., and Tuecke, S. 1996b. Software infrastructure for high-performance distributed computing. Preprint, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Ill.
- [9] Gropp, W., Lusk, E., Doss, N., and Skjellum, A. 1996. A high-performance, portable implementation of the MPI Message Passing Interface standard. Preprint MCS-P567-0296, Mathematics and Computer Science Division, Argonne National Laboratory, Argonne, Ill.
- [10] Korab, H., and Brown, M. D. 1995. Introduction. In Virtual environments and distributed computing at SC'95, New York and Los Alamitos, ACM/IEEE, pp. 1–6.
- [11] Partridge, C. 1994. Gigabit networking. Reading, Mass., Addison-Wesley.
- [12] Vetter, R. J. 1995. ATM concepts, architectures, and protocols. Communications of the ACM 38(2):30–38.